

ATLAS introduction

The Large Hadron Collider (LHC) at CERN will open a new frontier in particle physics due to its higher collision energy and luminosity compared to existing accelerators. The guiding principle in designing the ATLAS experiment, one of the two major LHC experiments, has been maximizing the discovery potential for new physics such as Higgs bosons and supersymmetric particles, while keeping the capability of high-accuracy measurements of known objects such as heavy quarks and gauge bosons. ATLAS is being constructed by 1850 collaborators in 150 institutes around the world. The detector employs precision tracking, calorimetry and muon measurement over a large solid angle to accurately identify and measure electrons, muons, photons, jets and missing energy. The very high interaction rate of 10^{34} Hz is filtered by a multiple trigger levels down to a recorded rate containing the rare interactions of interest of approximately 100 Hz. This corresponds to 100 MB/sec of raw data and 10 PB/year total raw and processed data volume. The scale of the computing challenge presented by the LHC experiments is unprecedented in terms of data volume, processing requirements, and the scale and distributed nature of the analysis task among thousands of physicists worldwide.

Distributed computing in ATLAS

Grid technologies are essential to fully realizing the potential of the ATLAS research program by enabling a collaboration-wide computing fabric that can deliver the capability for full participation in the development and execution of the ATLAS research program on the part of physicists at their home institutes. This is particularly true for distant participants such as those in the U.S. Meeting the demands of LHC analysis via a highly distributed, hierarchical computing infrastructure exploiting Grid technologies is a central element of the ATLAS worldwide computing model.

While LHC data taking will not begin until 2006, ATLAS already has a large and highly distributed computing and software operation serving immediate and near term needs such as test beam data analysis, detector performance and physics studies supporting detector design and optimization, software development and associated scalability studies, and 'Data Challenges' involving high-throughput, high-volume stress tests of offline processing software and facilities.

ATLAS is currently transitioning from FORTRAN-based software used for past production operations to new object oriented C++ software in which the U.S. has a leading role in architectural design and infrastructure development. In databases and data management, again with a leading U.S. presence, ATLAS has also recently begun ramping a development program to build the full ATLAS system. The U.S. has selected these focus areas as those best matched to U.S. expertise and to most effectively enabling U.S. physics analysis participation. Our PPDG program will provide an important complement to these efforts in providing a powerful and well integrated distributed computing capability that will empower physicists in performing analysis work at their home institutes.

PPDG Deliverables

Deliverables are roughly time-ordered throughout.

Year 1: Production distributed data service

The distributed data services described here are to exist between CERN, the Tier 1 Facility at BNL, and a few other U.S. institutes. Likely early participants are ANL, LBNL, Boston U, Indiana U, and U Michigan. The objective is a multi-point U.S. grid (in addition to the CERN link) providing distributed data services as early as possible.

- Distributed file and replica catalogs
 - Cataloging files resident on disk and in mass storage. Based on logical names in an agreed global Grid namespace. Serving read-only files at CERN, the Tier 1 at BNL, and select U.S. institutes. Supported by web, shell and API tools to browse, query and manage catalogs.
- Deployment of data transport services in production, serving the distributed data service
- Production distributed data service providing managed distribution and replication of data files

- User initiates data replication via web or command line. Simple replica cache management. Service is integrated with hierarchical mass storage: HPSS system at the BNL Tier 1; collaboration with CERN and EU DataGrid for CERN mass store support.
- Remote job submission service testbed (in use by developers and 'friendly users')
 - Supports remote submission of jobs to the Tier 1 from a number of external centers. Requires limited deployment of distributed authentication services.
- Grid instrumentation service: basic monitoring of distributed data service
 - Monitors activity, usage, performance, error conditions. Web based display and reporting tools.
- ATLAS 'data signature' design supporting Grid requirements in place. Prototyping begun.
 - A specification of data set content and characteristics complete enough to, in principle, regenerate it. Used in tests of dataset equivalence, current validity, etc.
- U.S. ATLAS distributed computing services architecture: requirements gathered, design begun.

Year 2: Production distributed job submission service

The major new functionality to be delivered in year 2 is a distributed job submission service.

- Broad deployment of user-level Grid authentication in support of 'Grid user' based functionality of production remote job submission service and enhanced distributed data service.
- Distributed data service enhancements
 - Extend service to several additional U.S. institutes. Integrate replica management into core database software supporting production. Add replica cache management service. Incorporate user-level storage resource reservation. Add cost estimation service. Provide C++ and Java APIs for catalog services. Automated, event-driven (eg. 'update available' signal) replica updating implemented to ensure consistency across file instances.
- Production remote job submission service
 - Extension of simple year 1 system to production use by the general community; requires full deployment of user-level Grid authentication.
- Extension of grid instrumentation service
 - Full monitoring of distributed data service and basic monitoring of job submission service. Improved display and reporting tools.

Year 3: Transparent distributed processing services

The principal goal in year 3 is to enhance the distributed data and processing services with user interfaces, specification languages, and ATLAS infrastructure integration to provide transparency to the user of the distributed nature of processing and analysis, both in 'batch' and interactively.

- Production distributed processing service
 - Extension of remote job submission service to provide transparency (to the distributed nature of the processing) to the ATLAS offline analysis user. Distributed services integrated directly into ATLAS software infrastructure.
- Integrated distributed data management services
 - Integration of distributed data services with ATLAS database and data management infrastructure to provide (policy-constrained) user transparency to data locality. Cost estimators supporting policy control integrated. Run, event, and event feature (tag) metadata integrated with PPDG catalogs.
- ATLAS 'data signature' deployed in support of coherence/consistency of file replicas and transparency in data set requests.

Associated major ATLAS milestones

Calorimeter test beam analysis: Summer/Fall 2001

Use as developmental testbed and early application area for distributed data service, making test beam data available at U.S. institutes involved in test beam analysis (ANL, BNL, possibly others)

ATLAS physics workshop: Sep 2001

Use as developmental testbed and early application area for distributed data service, making simulation data (produced for Physics TDR) available at U.S. institutes involved in physics studies (many U.S. institutes)

Data Challenge 1: Feb - Jul 2002

Major test at ~.1% scale relative to final system of data processing and computing facilities. Will incorporate distributed services then available in the Challenge.

Computing TDR preparation: May - Nov 2002

Experience to date with design and deployment of distributed data and processing services will drive the worldwide computing model developed for the Computing TDR.

Trigger/Data Acquisition TDR: Summer 2002

Heavy user of distributed data service.

Data Challenge 2: Jan - Sep 2003

As DC1, but at ~10% scale relative to final system, and including first large scale production deployment of multi-tier distributed computing services.

Physics Readiness Report: Jan – Jun 2004

Major report based on extensive physics studies. Full deployment of distributed services to participating physicists.

Full chain test: Jul 2004

Test of full processing chain at full bandwidth, from high level trigger through analysis. High throughput testing of distributed services.

20% Processing farm prototype: Dec 2004

Production processing test with 100% complexity (processor count), 20% capacity system relative to 2007 level. High throughput, high complexity testing of distributed services.